arXiv:1701.01370v1 [cs.CV] 5 Jan 2017

Learning from Synthetic Humans

Gül Varol*† Javier Romero‡ Xavier Martin† Naureen Mahmood‡

Michael Black‡ Ivan Laptev* Cordelia Schmid†

Abstract

Estimating human pose, shape, and motion from images and videos are fundamental challenges with many applications. Recent advances in 2D human pose estimation use large amounts of manually-labeled training data for learning convolutional neural networks (CNNs). Such data is time consuming to acquire and difficult to extend. Moreover, manual labeling of 3D pose, depth and motion is impractical. In this work we present SURREAL (Synthetic hUmans foR REAL tasks): a new large-scale dataset with synthetically-generated but realistic images of people rendered from 3D sequences of human motion capture data. We generate more than 6 million frames together with ground truth pose, depth maps, and segmentation masks. We show that CNNs trained on our synthetic dataset allow for accurate human depth estimation and human part segmentation in real RGB images. Our results and the new dataset open up new possibilities for advancing person analysis using cheap and large-scale synthetic data.

1. Introduction

Convolutional Neural Networks provide significant gains to problems with large amounts of training data. In the field of human analysis, recent datasets [4, 32] now gather a sufficient number of annotated images to train networks for 2D human pose estimation [21, 34]. Other tasks such as accurate estimation of human motion, depth and body-part segmentation are lagging behind as manual supervision for such problems at large scale is prohibitively expensive.

Images of people have rich variation in poses, clothing, hair styles, body shapes, occlusions, viewpoints, motion blur and other factors. Much of these variations, however, can be synthesized using existing 3D motion capture (MoCap) data [3, 16] and modern tools for realistic rendering. Provided sufficient realism, such an approach would be highly useful for many tasks as it can generate rich ground truth in terms of depth, motion, body-part segmentation and occlusions to mention a few.

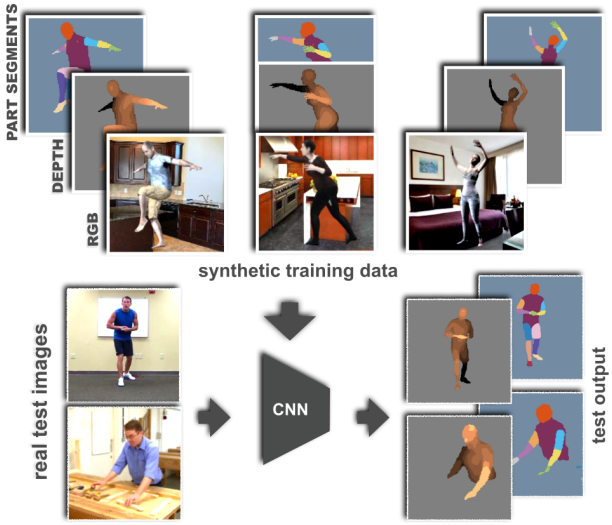

Figure 1. We generate photo-realistic synthetic images and their corresponding ground truth for learning pixel-wise classification problems: human parts segmentation and depth estimation. The convolutional neural network trained only on synthetic data generalizes on real images sufficiently for both tasks. Real test images in this figure are taken from MPII Human Pose dataset [4].

In this work we present SURREAL: a new large-scale dataset with synthetically-generated but realistic images of people. Images are rendered from 3D sequences of MoCap data. To ensure realism, the synthetic bodies are created using the SMPL body model [18], whose parameters are fit by the MoSh [19] method given raw 3D MoCap marker data. We randomly sample a large variety of viewpoints, clothing and lighting. SURREAL contains more than 6 million frames together with ground truth pose, depth maps, and segmentation masks. We show that CNNs trained on synthetic data allow for accurate human depth estimation and human part segmentation in real RGB images, see Figure 1. SURREAL dataset will become publicly available [1].

The rest of this paper is organized as follows. Section 2

*WILLOW project-team, Département d'Informatique de l'École Normale Supérieure, ENS/Inria/CNRS UMR 8548, Paris, France.

†Thoth project-team, Inria Grenoble Rhône-Alpes, Laboratoire Jean Kuntzmann, France.

‡Max Planck Institute for Intelligent Systems, Tübingen, Germany.

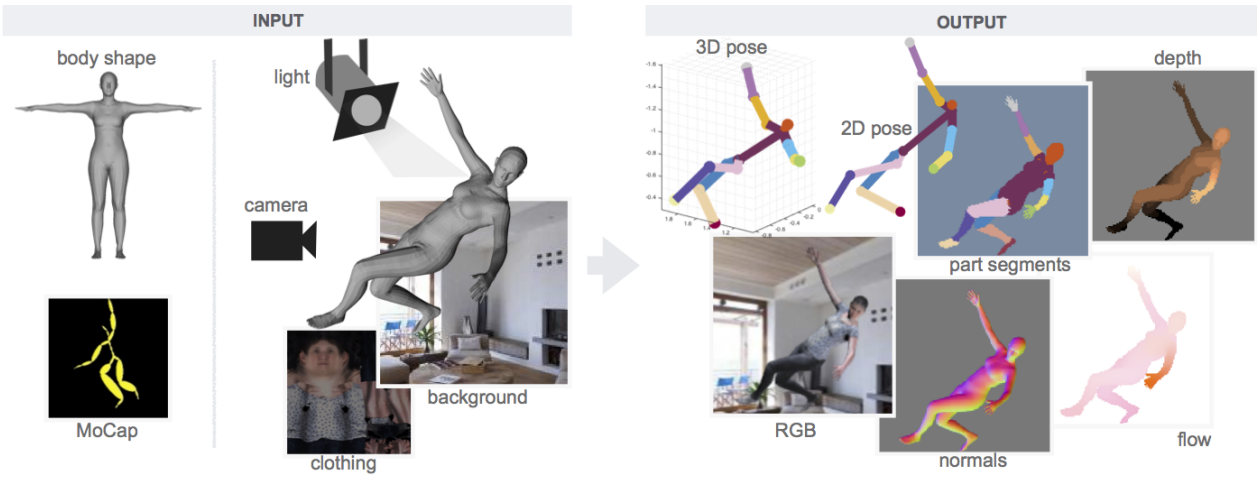

Figure 2. Our pipeline for generating synthetic data. A 3D human body model is posed using motion capture data and a frame is rendered using a background image, a texture map on the body, lighting and a camera position. These ingredients are randomly sampled to increase the diversity of the data. We generate RGB images together with 2D/3D poses, surface normals, optical flow, depth images, and body-part segmentation maps for rendered people.

reviews previous work on the use of synthetic datasets in computer vision. Section 3 presents our approach for generating realistic synthetic videos of people. In Section 4 we describe our convolution architecture for human body-part segmentation and depth estimation. Section 5 reports experiments. We conclude in Section 6.

2. Related work

Knowledge transfer from synthetic to real images has been recently studied with deep neural networks. Dosovitskiy *et al*. [8] learn a CNN for optical flow estimation using synthetically generated images of rendered 3D moving chairs. Peng *et al*. [23] study the effect of different visual cues such as object/background texture and color when rendering synthetic 3D objects for object detection task. Similarly, [33] explores rendering 3D objects to perform viewpoint estimation. Recently, Gaidon *et al*. [12] have released the Virtual KITTI dataset with synthetically generated videos of cars to study multi-object tracking.

Several works focused on creating synthetic images of humans for learning 2D pose estimation [24, 31], 3D pose estimation [7, 9, 13, 30, 37], pedestrian detection [20, 24, 25], and action recognition [27, 28]. Pishchulin *et al*. [25] generate synthetic images with a game engine. In [24], they deform 2D images with a 3D model. More recently, Rogez and Schmid [30] used image-based synthesis engine to augment existing real images. Ghezelghieh *et al*. [13] render synthetic images with 10 simple body models with an emphasis on upright people; however, the main challenge using existing MoCap data for training is to generalize to poses that are not upright.

A similar direction has been explored in [27, 28]. In [27], action recognition problem is addressed with synthetic human trajectories from MoCap data. [28] trains CNNs with synthetic depth images.

The closest work to this paper is [7], where the authors render large-scale synthetic images for predicting 3D pose with CNNs. Our dataset differs from [7] by having a richer, per-pixel ground truth, thus allowing to train for pixel-wise predictions and multi-task scenarios. In addition, we argue that the realism in our synthetic images is better (see [1]), thus resulting in a smaller gap between features learned from synthetic and real images. The method in [7] heavily relies on real images as input in their training with domain adaptation. This is not the case for our synthetic training. Moreover, we render video sequences which can be used for temporal modeling.

In this paper, we show that photo-realistic renderings of people under large variations in shape, texture, viewpoint and pose can help to solve pixel-wise human labeling tasks. We use synthesized data to learn models of people and demonstrate improvements of body part segmentation and body depth estimation in real images.

3. Data generation

This section presents our SURREAL (Synthetic hUmans foR REAL tasks) dataset and describes key steps for its generation (Section 3.1). We also describe how we obtain ground truth data for real MoCap sequences (Section 3.2).

3.1. Synthetic humans

Our pipeline for generating synthetic data is illustrated in Figure 2. A human body with a *random* 3D pose, *random* shape and *random* texture is rendered from a *random* viewpoint for some *random* lighting and a *random* background image. Below we define what "random" means in all these

Figure 3. Sample synthetic frames from our SURREAL dataset. Images contain realistically rendered people with a large variety of poses, shapes, clothings, viewpoints and background clutter.

cases. Since the data is synthetic, we also generate ground truth depth maps, optical flow, surface normals, human part segmentations and joint locations (both 2D and 3D). As result, we obtain 6.5 million frames grouped into $67,582$ continuous image sequences. See Table 1 for more statistics and Figure 3 for samples from the SURREAL dataset.

Body model. Synthetic bodies are created using the SMPL body model [18]. SMPL is a realistic articulated model of the body created from thousands of high-quality 3D scans, which decomposes body deformations into pose (kinematic deformations due to skeletal posture) and shape (body deformations intrinsic to a particular person that make them different from others). SMPL is compatible with most animation packages like Blender [2]. SMPL deformations are modeled as a combination of linear blend skinning and linear blendshapes defined by principal components of body shape variation. SMPL pose and shape parameters are converted to a triangulated mesh using Blender, which then applies texture, shading and adds a background to generate the final RGB output.

Body shape. In order to render varied, but realistic, body shapes we make use of the CAESAR dataset [29], which was used to train SMPL. To create a body shape, we select one of the CAESAR subjects at random and approximate their shape with the first 10 SMPL shape principal components. Ten shape components explain more than 95% of the shape variance in CAESAR (at the resolution of our mesh) and produce quite realistic body shapes.

Body pose. To generate images of people in realistic poses, we take motion capture data from the CMU MoCap database [3]. CMU MoCap contains more than 2000 sequences of 23 high-level action categories, resulting in more than 10 hours of recorded 3D locations of body markers.

It is often challenging to realistically and automatically retarget MoCap skeleton data to a new model. Consequently we do not use the skeleton data but rather use MoSh [19] to fit the SMPL parameters that best explain raw 3D MoCap marker locations. This gives both the 3D shape of the subject and the articulated pose parameters of SMPL. To increase the diversity, we replace the estimated 3D body shape with a set of randomly sampled body shapes.

We render each CMU MoCap sequence three times using different random view-points. Moreover, we divide the sequences into chunks of 100 frames with 30%, 50% and 70% overlaps for these three renderings. Every pose of the sequence is rendered with consistent parameters (i.e. body shape, clothing, light, background etc.) within each 100-frame chunk.

Human texture. We texture the body model with texture maps extracted from two types of real scans. The first one uses CAESAR scans, which come with a color texture per 3D point, and we extract SMPL texture maps from the scans. These texture maps are varied in terms of skin color and identity. However, their quality is not ideal due to low resolution, uniform tight-fitting clothing, and visible markers placed on the face and body. Anthropometric markers are automatically removed from the texture images and inpainted.

In order to provide more variety, we extract a second set of textures from 3D scans of 7 different subjects with normal clothing from a total of 98 different capture sessions. These scans are registered with 4Cap as in [26]. The texture of real clothing substantially increases the realism of generated images, even though SMPL does not model 3D deformations of clothes.

20% of our data is rendered with the first set (208 CAESAR textures randomly sampled from 4000), and the rest with the second set (701 clothed textures). To preserve the anonymity of subjects, we replace all faces in the texture

maps by the average CAESAR face. The skin color of this average face is corrected to fit the face skin color of the original texture map. This corrected average face is blended smoothly with the original map, resulting in a realistic and anonymized body texture.

Light. The body is illuminated using Spherical Harmonics with 9 coefficients [14]. The coefficients are randomly sampled from a uniform distribution between -0.7 and 0.7, apart from the ambient illumination coefficient (which has a minimum value of 0.5) and the vertical illumination component, which is biased to encourage the illumination from above. Since Blender does not provide Spherical Harmonics illumination, a spherical harmonic shader for the body material was implemented in Open Shading Language (OSL).

Camera. The projective camera has a resolution of 320×240, focal length of 60mm and sensor size of 32mm. To generate images of the body in a wide range of positions, we take 100-frame MoCap sub-sequences and, in the first frame, render the body so that the center of the viewport points to the pelvis of the body, at a random distance (sampled from a normal distribution with 8 meters mean, 1 meter deviation) with a random yaw angle. The remainder of the sequence then effectively produces bodies in a range of locations relative to the camera.

Background. We render the person on top of a static background image. To ensure that the backgrounds are reasonably realistic and do not include other people, we sample from a subset of LSUN dataset [36] that includes total of 400K images from the categories kitchen, living room, bedroom and dining room.

Ground truth. We perform multiple rendering passes in Blender to generate different types of per-pixel ground truth. The *material* pass generates pixel-wise segmentation of rendered body parts, given different material indices assigned to different parts of our body model. The *velocity* pass, typically used to simulate motion blur, provides us with a render simulating ground truth optical flow. The *depth* and *normal* passes, used for emulating effects like fog, bokeh or for performing shading, produces per-pixel depth maps and normal maps. The final texture rendering pass composites the shaded, textured body over the random background. Together with this data we also save camera parameters, lighting parameters as well as the 2D and 3D positions of body joints.

3.2. Generating ground truth for real human data

Human3.6M dataset [16, 15] provides ground truth for 2D and 3D human poses. We complement this ground truth and generate predicted body-part segmentation and depth maps for people in Human3.6M. Here again we use MoSh [19] to fit the SMPL body shape and pose to the raw MoCap marker data. This provides a good fit of the model to the shape and the pose of real bodies. Given the provided camera calibration, we project models to images. We then render the ground truth segmentation, depth, and 2D/3D joints just as above. The only difference here is that these now are in correspondence with real pixel values in the dataset. As MoSh provides almost perfect fits of the model, we consider this data to be "ground truth". See Figures 6 and 7 for generated examples.

Table 1. SURREAL dataset in numbers. Each MoCap sequence is rendered 3 times (with 3 different overlap ratios). Clips are mostly 100 frames long. We obtain a total of 6,5million frames.

	#subjects	#sequences	#clips	#frames
Train	115	1,964	55,001	5,342,090
Test	30	703	12,528	1,194,662
Total	145	2,607	67,582	**6,536,752**

4. Approach

In this section, we present our approach for human body part segmentation [22, 5] and human depth estimation [10, 11, 17], which we train with synthetic and/or real data, see Section 5 for an evaluation.

Our approach builds on the stacked hourglass network architecture introduced originally for 2D pose estimation problem [21]. This network involves several repetitions of contraction followed by expansion layers which have skip connections to implicitly model spatial relations from different resolutions that allows bottom-up and top-down structured prediction. The convolutional layers with residual connections and 8 'hourglass' modules are stacked on top of each other, each successive stack taking the previous stack's prediction as input. The reader is referred to [21] for more details. A variant of this network is used for scene depth estimation [6]. We choose this architecture because it can infer pixel-wise output by taking into account human body structure.

Our network input is a 3-channel RGB image of size 256×256 cropped and scaled to fit a human bounding box using the ground truth. The network output for each stack has dimensions $64 \times 64 \times 15$ in the case of segmentation (14 classes plus the background) and $64 \times 64 \times 20$ for depth (19 depth classes plus the background). We use cross-entropy loss defined on all pixels for both segmentation and depth. The final loss of the network is the sum over 8 stacks. We train for 50K iterations for synthetic pre-training using the RMSprop algorithm with mini-batches of size 6 and a learning rate of 10^{-3}. Our data augmentation during training includes random rotations, scaling and color jittering.

We formulate the problem as pixel-wise classification task for both segmentation and depth. When addressing

segmentation, each pixel is assigned to one of the predefined 14 human parts, namely head, torso, upper legs, lower legs, upper arms, lower arms, hands, feet (separately for right and left) or to the background class. Regarding the depth, we align ground-truth depth maps on the z-axis by the depth of the pelvis joint, and then quantize depth values into 19 bins (9 behind and 9 in front of the pelvis) We set the quantization constant to 45mm to roughly cover the depth extent of common human poses. The network is trained to classify each pixel into one of the 19 depth bins or background. At test time, we first upsample feature maps of each class with bilinear interpolation by a factor of 4 to output the original resolution. Then, each pixel is assigned to the class for which the corresponding channel has the maximum activation.

5. Experiments

We test our approach on several datasets. First, we validate whether our models for segmentation and depth estimation perform well on the synthetic test set of our SURREAL dataset. This gives an idea about the model's ability to learn the given task for the same domain. Second, we evaluate the segmentation task on real images from the Freiburg Sitting People dataset [22]. Next, we measure segmentation and depth estimation performance on videos from the Human3.6M dataset [16, 15], where the 3D information is also available. Finally, we qualitatively evaluate our approach on the MPII Human Pose dataset [4] to see success and failure cases in a challenging domain adaptation scenario.

5.1. Evaluation metrics

We choose intersection over union (IOU) and pixel accuracy metrics for evaluating the segmentation approach. The final measure is the average over 14 human parts as in [22]. Depth estimation is formulated as a classification problem, but we take into account the continuity when we evaluate. We use root-mean-squared-error (RMSE) that is computed between the predicted quantized depth value (class) and the ground truth quantized depth on the human pixels. However, to be able to interpret the error in terms of real world coordinates, we multiply the error by the quantization constant that we used, which is 45mm. Furthermore, we report a scale and translation invariant RMSE (st-RMSE) by solving for the best translation and scaling in z-axis to fit the prediction to the ground truth. Since inferring depth from RGB is ambiguous, this is a common technique used in evaluations [11].

5.2. Validation on synthetic images

Train/test split. To evaluate our methods on synthetic images, we separate 20% of the synthetic frames for the test set and train all our networks on the remaining training set.

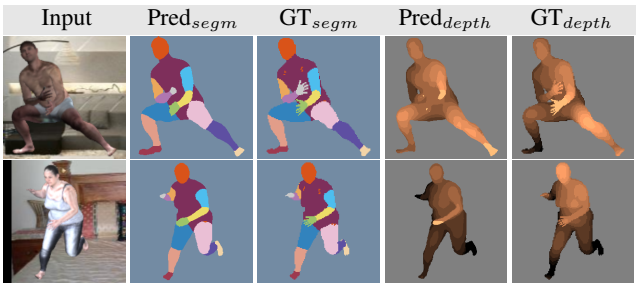

Figure 4. Segmentation and depth predictions on our synthetic test set.

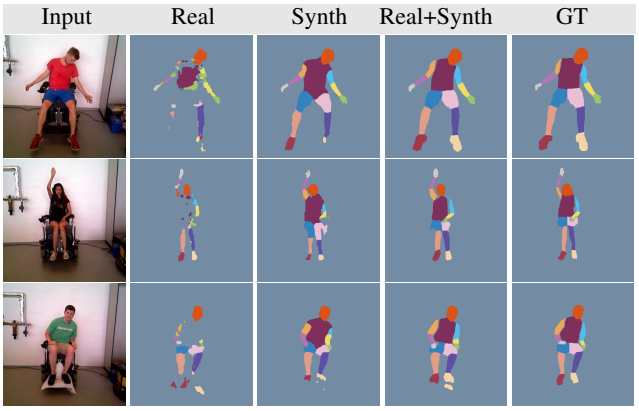

Figure 5. Part segmentation on the Freiburg Sitting People dataset, training only on FSitting (Real), training only on synthetic images (Synth), fine-tuning on 2 training subjects from FSitting (Real+Synth). Fine-tuning helps although only for 200 iterations.

The split is constructed such that a given CMU MoCap subject is assigned as either train or test. However some subjects have a large number of instances, some subjects have unique actions, and some actions are very common (walk, run, jump). Overall, 30 subjects out of 145 are assigned as test. 28 test subjects cover all common actions, and 2 test subjects have unique actions. Remaining subjects are used for training. Although our synthetic images have different body shape and appearance than the subject in the originating MoCap sequence, we still found it appropriate to split by subjects. We separate a subset of our body shapes, clothing and background images for the test set. This ensures that our tests are unbiased with regards to appearance, yet are still representative of all actions. Table 1 summarizes the number of frames, clips and MoCap sequences in each split. Clips are the continuous 100-frame image sequences where we have the same random body shape, background, clothing, camera and lighting. A new random set is picked at every clip. Note that a few sequences have less than 100 frames, therefore not all the clips have 100-frames.

Results on synthetic test set. The evaluation is performed on the middle frame of each 100-frame clip on the aforementioned held-out synthetic test set, totaling in

Table 2. Parts segmentation results on 4 test subjects of Freiburg Sitting People dataset. IOU for head, torso and upper legs (averaged over left and right) are presented as well as the mean IOU and mean pixel accuracy over 14 parts. The means do not include background class. By adding an upsampling layer, we get the best results reported on this dataset.

Training data	Head IOU	Torso IOU	Legs_{up} IOU	mean IOU	mean Acc.
Real+Pascal[22]	-	-	-	64.10	81.78
Real	58.44	24.92	30.15	28.77	38.02
Synth	73.20	65.55	39.41	40.10	51.88
Synth+Real	72.88	80.76	65.41	59.58	78.14
Synth+Real+up	**85.09**	**87.91**	**77.00**	**68.84**	**83.37**

Table 3. Parts segmentation results on Human3.6M. The best result is obtained by fine-tuning synthetic network with real images. Although the performance of the network trained only with real data outperforms training only with synthetic, the predictions visually are worse because of overfitting, see Figure 6.

	IOU		Accuracy	
Training data	fg+bg	fg	fg+bg	fg
Real	49.61	46.32	58.54	55.69
Synthetic	46.35	42.91	56.51	53.55
Synthetic+Real	57.07	54.30	67.72	65.53

12,528 images. For segmentation, the IOU and pixel accuracy are 69.13% and 80.61%, respectively. Evaluation of depth estimation gives 72.9mm and 56.3mm for RMSE and st-RMSE errors, respectively. Figure 4 shows sample predictions. For both tasks, the results are mostly accurate on synthetic test images. However, there exist a few challenging poses (e.g. crawling), test samples with extreme close-up views, and fine details of the hands that are causing errors. In the following sections, we investigate if similar conclusions can be made for real images.

5.3. Segmentation on Freiburg Sitting People

Freiburg Sitting People (FSitting) dataset [22] is composed of 200 high resolution (300x300 pixels) front view images of 6 subjects sitting on a wheel chair. There are 14 human part annotations available. See Figure 5 for sample test images and corresponding ground truth (GT) annotation. We use the same train/test split as [22], 2 subjects for training and 4 subjects for test. The amount of data is limited for training deep networks. We show that our network pre-trained only on synthetic images is already able to segment human body parts. This shows that the human renderings in the synthetic dataset are representative of the real images, such that networks trained exclusively on synthetic data can generalize quite well to real data.

Table 2 summarizes segmentation results on FSitting. We carry out several experiments to understand the gain from synthetic pre-training. For the 'Real' baseline, we train the same network from scratch using 2 training subjects. This network quickly overfits as there are few subjects to learn from and the performance is quite low. Our 'Synthetic' result is obtained using the network pre-trained on synthetic images without fine-tuning. We get 51.88% pixel accuracy and 40.1% IOU with this method and clearly outperform training from real images. Furthermore, fine-tuning ('Synthetic+Real') with 2 training subjects helps significantly. See Figure 5 for qualitative results. Given the little amount for training in FSitting, the fine-tuning quickly converges after 200 iterations.

In [22], the authors introduce a network that outputs a high-resolution segmentation after several layers of upconvolutions. For a fair comparison, we modify our network to output full resolution by adding one bilinear upsampling layer followed by nonlinearity (ReLU) and a convolutional layer with 3×3 filters that outputs $15 \times 300 \times 300$ instead of $15 \times 64 \times 64$ as explained in Section 4. If we fine-tune this network ('Synth+Real+up') on FSitting, we improve performance and outperform [22] by a large margin. Note that [22] trains on the same FSitting training images, but added around 2,800 Pascal images. Hence they use significantly more manual annotation than our method.

5.4. Segmentation and depth on Human3.6M

Showcasing our approach requires expensive annotation if applied on real data. This is the reason for existing real datasets being small and not suitable for training deep networks. As a result, there is no available dataset to carry out a clean evaluation. For this reason, we came up with generating nearly perfect ground truth for images recorded with a calibrated camera and given their MoCap data. Human3.6M is currently the largest dataset where such information is available. There are 3.6 million frames from 4 cameras. We use subjects 1, 5, 6, 7, 8 for training, subject 9 for validation and subject 11 for testing as in [30, 35]. We use all frames from one repetition of each action for training, and every 64^{th} frame from all repetitions at test. The frames have resolution 1000×1000 pixels, we assume a 256×256 cropped human bounding box is given to reduce computational complexity. We evaluate the performance of both segmentation and depth, and compare with the baseline for which we train a network on real images only.

5.4.1 Segmentation

Table 3 summarizes the parts segmentation results on Human3.6M. We report both the mean over 14 human parts (fg) and the mean together with the background class (fg+bg). Training on real images instead of synthetic images increases IOU by 3.4% and pixel accuracy by 2.14%. This is expected because the training distribution matches the test distribution in terms of background, camera position and action categories (i.e. poses). The amount of real

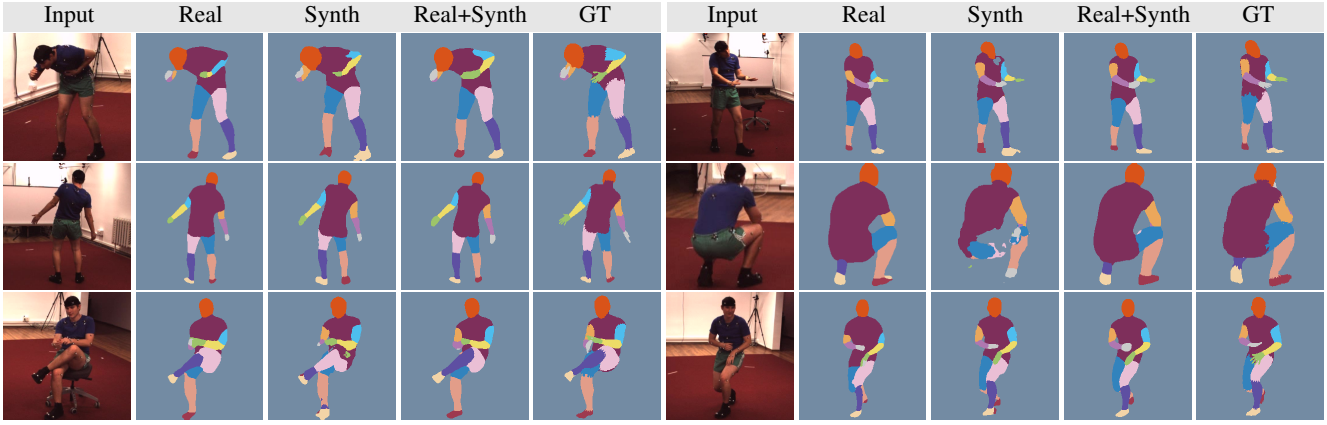

Figure 6. Parts segmentation on the Human3.6M dataset, training only on real images and MoSH-generated ground-truth from Human3.6M (Real), training only on synthetic images from SURREAL (Synth), and fine-tuning on real Human3.6M data (Real+Synth). The 'Real' baseline clearly fails on upper arms by fitting the skin color. The synthetic pre-trained network has seen more variety in clothing. Best result is achieved by the fine-tuned network.

data is sufficient to perform CNN training. However, since there are very few subjects available, we see that the network doesn't generalize to different clothing. In Figure 6, the 'Real' baseline has the border between shoulders and upper arms exactly on the t-shirt boundaries. This reveals that the network learns about skin color rather than actual body parts. Our pre-trained network (Synth) performs reasonably well, even though the pose distribution in our MoCap is quite different than that of Human3.6M. When we fine-tune the network with real images from Human3.6M (Real+Synth), the segmentations look very similar to the ground truth and we outperform the baseline by a large margin. Moreover, our model is capable of distinguishing left and right most of the time on all 4 views since it has been trained with randomly sampled views.

5.4.2 Depth estimation

Depth estimation results on Human3.6M for various poses and viewpoints are illustrated in Figure 7. Here, the pre-trained network fails at the very challenging poses, although it still captures partly correct estimates (left column, third, fourth rows). Fine-tuning on real data compensates for these errors and refines estimations. In Table 4, we show RMSE error measured on foreground pixels, together with the scale-translation invariant version (see Section 5.1). We also report the error only on known 2D joints (PoseRMSE) to have an idea of how well a 3D pose estimation model would work based on the depth predictions. One would need to handle occluded joints to infer 3D locations of all joints, and this is beyond the scope of the current paper.

5.5. Qualitative results on MPII Human Pose

The datasets we evaluated our approach with are relatively simple, i.e. no background clutter, few subjects, single person per image, full body visible. Except the first two properties, this setup matches our synthetic training data.

Table 4. Depth estimation results on Human3.6M (in millimeters). The depth error (RMSE) and scale-translation invariant error (st-RMSE) are reported on foreground pixels. The error measurement only on given human joints is also provided (PoseRMSE).

Training data	RMSE	st-RMSE	PoseRMSE	st-PoseRMSE
Real	96.3	75.2	122.6	94.5
Synthetic	111.6	98.1	152.5	131.5
Synthetic+Real	90.0	67.1	92.9	82.8

Therefore, we test the generalization of our model on more challenging images. Currently, MPII Human Pose [4] is one of the biggest datasets with diverse viewpoints and clutter. However, this dataset has no annotation for part segmentation nor depth. For this reason, we qualitatively show our predictions. In Figure 8, several success and failure cases are illustrated. Our model generalizes reasonably well, except when there are multiple people close to each other, and in extreme viewpoints and occlusions, which have not appeared during training.

6. Conclusions

In this study, we have shown successful large-scale training of CNNs from synthetically generated images or people. We have addressed two tasks, namely, human body parts segmentation and depth estimation, for which large-scale manual annotation is infeasible. Our generated synthetic dataset comes with rich pixel-wise ground truth information and can potentially be used for other tasks than considered here. Unlike many existing synthetic datasets, the focus of SURREAL is on the realistic rendering of people, which is a challenging task. In our future work, we plan to integrate the person into the background in a more realistic way by taking into account the lighting and the 3D scene layout. We also plan to augment the data with more challenging scenarios by creating occlusions and rendering multiple people in the same scene.

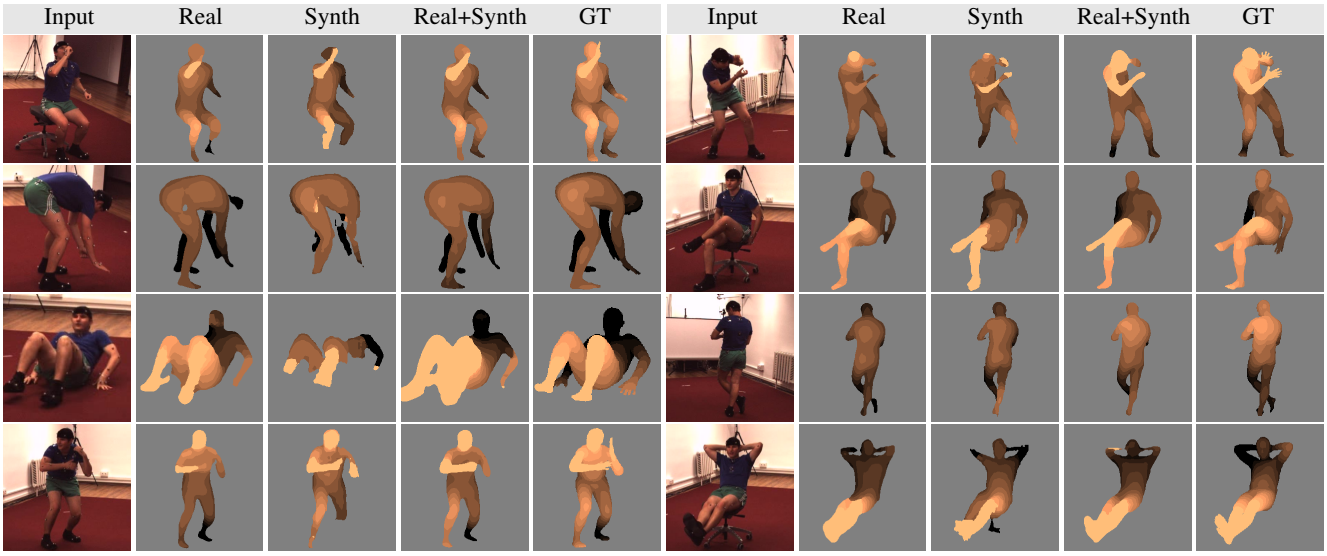

Figure 7. Depth segmentation on the Human3.6M dataset, columns represent same training partitions as in Figure 6. The pre-trained network (Synth) fails due to scale mismatching in the training set and low contrast body parts, but fine-tuning with real data (Real+Synth) tends to recover from these problems.

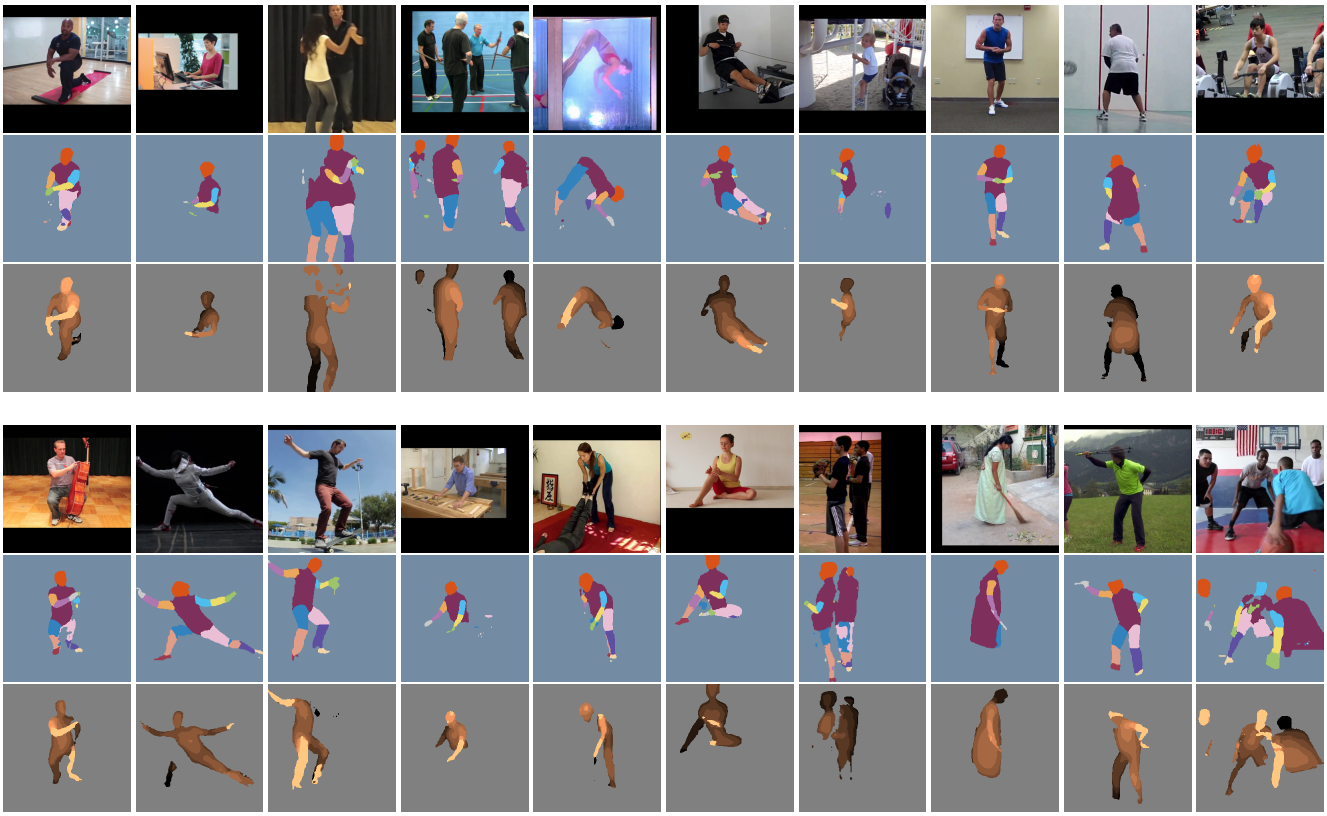

Figure 8. Qualitative results on challenging images from MPII Human Pose dataset. Multi-person, occlusion and extreme poses are difficult cases for our model. Given that there is no fine-tuning involved, the model is able to generalize sufficiently well on cluttered real data.

Acknowledgements

This work was supported by the MSR-Inria joint lab, Google and Facebook Research Awards, the ERC starting grant ACTIVIA and the ERC advanced grant ALLEGRO. We also acknowledge the funding provided by the Alexander von Humbolt Foundation.

References

[1] http://www.di.ens.fr/willow/research/surreal/. 1, 2

[2] Blender - a 3D modelling and rendering package. http://www.blender.org. 3

[3] Carnegie-Mellon Mocap Database. http://mocap.cs.cmu.edu/. 1, 3

[4] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 1, 5, 7

[5] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *CVPR*, 2016. 4

[6] W. Chen, Z. Fu, D. Yang, and J. Deng. Single-image depth perception in the wild. In *NIPS*, 2016. 4

[7] W. Chen, H. Wang, Y. Li, H. Su, Z. Wang, C. Tu, D. Lischinski, D. Cohen-Or, and B. Chen. Synthesizing training images for boosting human 3D pose estimation. In *3DV*, 2016. 2

[8] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 2

[9] Y. Du, Y. Wong, Y. Liu, F. Han, Y. Gui, Z. Wang, M. Kankanhalli, and W. Geng. Marker-less 3D human motion capture with monocular image sequence and height-maps. In *ECCV*, 2016. 2

[10] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. 4

[11] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014. 4, 5

[12] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR*, 2016. 2

[13] M. F. Ghezelghieh, R. Kasturi, and S. Sarkar. Learning camera viewpoint using cnn to improve 3D body pose estimation. In *3DV*, 2016. 2

[14] R. Green. Spherical harmonic lighting: The gritty details. In *Archives of the Game Developers Conference*, volume 56, 2003. 4

[15] C. Ionescu, L. Fuxin, and C. Sminchisescu. Latent structured models for human pose estimation. In *ICCV*, 2011. 4, 5

[16] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 1, 4, 5

[17] F. Liu, C. Shen, and G. Lin. Deep convolutional neural fields for depth estimation from a single image. In *CVPR*, 2015. 4

[18] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *SIGGRAPH Asia*, 2015. 1, 3

[19] M. M. Loper, N. Mahmood, and M. J. Black. MoSh: Motion and shape capture from sparse markers. *SIGGRAPH Asia*, 2014. 1, 3, 4

[20] J. Marin, D. Vazquez, D. Geronimo, and A. M. Lopez. Learning appearance in virtual scenarios for pedestrian detection. In *CVPR*, 2010. 2

[21] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016. 1, 4

[22] G. Oliveira, A. Valada, C. Bollen, W. Burgard, and T. Brox. Deep learning for human part discovery in images. In *ICRA*, 2016. 4, 5, 6

[23] X. Peng, B. Sun, K. Ali, and K. Saenko. Learning deep object detectors from 3D models. In *ICCV*, 2015. 2

[24] L. Pishchulin, A. Jain, M. Andriluka, T. Thormhlen, and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR*, 2012. 2

[25] L. Pishchulin, A. Jain, C. Wojek, M. Andriluka, T. Thormhlen, and B. Schiele. Learning people detection models from few training samples. In *CVPR*, 2011. 2

[26] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black. Dyna: A model of dynamic human shape in motion. *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, 2015. 3

[27] H. Rahmani and A. Mian. Learning a non-linear knowledge transfer model for cross-view action recognition. In *CVPR*, 2015. 2

[28] H. Rahmani and A. Mian. 3D action recognition from novel viewpoints. In *CVPR*, 2016. 2

[29] K. Robinette, S. Blackwell, H. Daanen, M. Boehmer, S. Fleming, T. Brill, D. Hoeferlin, and D. Burnsides. Civilian American and European Surface Anthropometry Resource (CAESAR), Final Report. 2002. 3

[30] G. Rogez and C. Schmid. MoCap-guided Data Augmentation for 3D Pose Estimation in the Wild. In *NIPS*, 2016. 2, 6

[31] J. Romero, M. Loper, and M. J. Black. FlowCap: 2D human pose from optical flow. In *GCPR*, 2015. 2

[32] B. Sapp and B. Taskar. Multimodal decomposable models for human pose estimation. In *CVPR*, 2013. 1

[33] H. Su, C. R. Qi, Y. Li, and L. J. Guibas. Render for CNN: Viewpoint estimation in images using CNNs trained with rendered 3D model views. In *ICCV*, 2015. 2

[34] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. *CVPR*, 2016. 1

[35] H. Yasin, U. Iqbal, B. Krger, A. Weber, and J. Gall. A dual-source approach for 3D pose estimation from a single image. In *CVPR*, 2016. 6

[36] F. Yu, Y. Zhang, S. Song, A. Seff, and J. Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv:1506.03365*, 2015. 4

[37] X. Zhou, M. Zhu, S. Leonardos, K. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3D human pose estimation from monocular video. In *CVPR*, 2016. 2